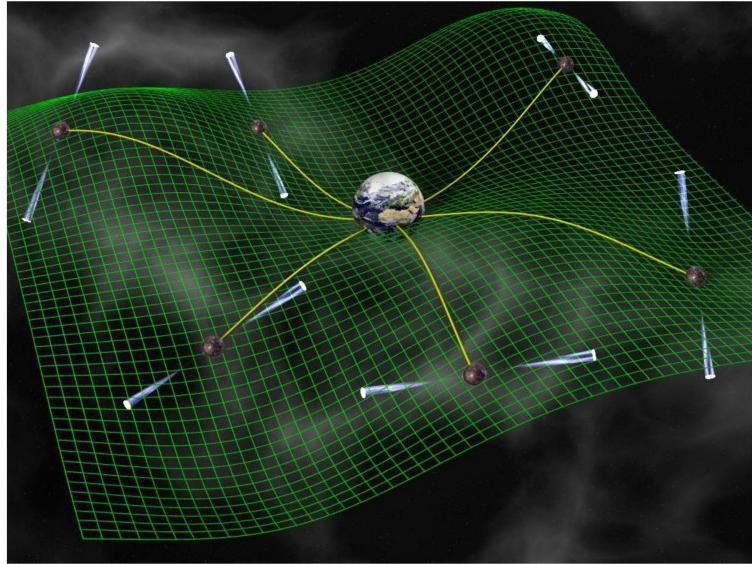


# The North American Nanohertz Observatory for Gravitational Waves

An Astro2010 decadal survey activity submission



**Principal Authors:** F. Jenet (Center for Gravitational Wave Astronomy @ UT Brownsville, 626-524-0518, [merlyn@phys.utb.edu](mailto:merlyn@phys.utb.edu)); L. S. Finn (Penn State, 814-863-9598, [lsfinn@psu.edu](mailto:lsfinn@psu.edu)); J. Lazio (NRL, 202-404-6329, [Joseph.Lazio@nrl.navy.mil](mailto:Joseph.Lazio@nrl.navy.mil)); A. Lommen (Franklin and Marshall College, 717-291-3810, [alommen@fandm.edu](mailto:alommen@fandm.edu)); M. McLaughlin (West Virginia University, 304-293-3422, [maura.mclaughlin@mail.wvu.edu](mailto:maura.mclaughlin@mail.wvu.edu)); I. Stairs (UBC, 604-822-6796, [stairs@astro.ubc.ca](mailto:stairs@astro.ubc.ca)); D. Stinebring (Oberlin, 440-775-8331, [dan.stinebring@oberlin.edu](mailto:dan.stinebring@oberlin.edu)); J. Verbist (WVU, 304-293-3422x1462, [Joris.Verbiest@mail.wvu.edu](mailto:Joris.Verbiest@mail.wvu.edu));

**NANOGrav Members and Contributors:** A. Archibald (McGill); Z. Arzoumanian (CRESST/USRA/NASA-GSFC); D. Backer (UC Berkeley); J. Cordes (Cornell); P. Demorest (NRAO); R. Ferdman (CNRS, France); P. Freire (NAIC); M. Gonzalez (UBC); V. Kaspi (McGill); V. Kondratiev (WVU); D. Lorimer (WVU); R. Lynch (Virginia); D. Nice (Bryn Mawr); S. Ransom (NRAO); R. Shannon (Cornell); X. Siemens (UW Milwaukee);

# Contents

<b>1</b>	<b>Summary</b>	<b>2</b>
<b>2</b>	<b>Key Science Goals</b>	<b>3</b>
2.1	Background . . . . .	3
2.2	Understand the co-evolution of galaxies and supermassive black holes . . . . .	4
2.3	Search for the signatures of early-universe or exotic physics processes . . . . .	5
2.4	Probe the nature of space-time . . . . .	6
2.5	Discover previously unanticipated gravitational wave sources . . . . .	7
<b>3</b>	<b>Technical Overview</b>	<b>8</b>
3.1	Observations . . . . .	9
3.2	Pulse Time of Arrival Determination . . . . .	10
3.3	Timing Model Analysis . . . . .	10
3.4	Correlation Analysis . . . . .	11
3.5	New Pulsars for the Timing Array . . . . .	12
3.6	Justification of Required Telescope Time . . . . .	12
<b>4</b>	<b>Technology Drivers</b>	<b>13</b>
4.1	Algorithm Development . . . . .	13
4.2	Removing the Effects of the Interstellar Medium . . . . .	14
4.3	Pulsar Timing with Interferometers . . . . .	15
<b>5</b>	<b>Activity Organization</b>	<b>16</b>
<b>6</b>	<b>Activity Schedule</b>	<b>17</b>
<b>7</b>	<b>Cost Estimates</b>	<b>19</b>

# 1 Summary

The North American Nanohertz Observatory for Gravitational Waves (NANOGrav) is a consortium of astronomers whose goal is the creation of a galactic scale gravitational wave observatory sensitive to gravitational waves in the nHz –  $\mu$ Hz band. It is just one component of an international collaboration involving similar organizations of European and Australian astronomers who share the same goal.

Gravitational waves, a prediction of Einstein’s general theory of relativity, are a phenomenon of dynamical space-time generated by the bulk motion of matter (e.g., the rapid periapsis passage of star on a low angular momentum orbit about a supermassive black hole), primordial or quantum fluctuations arising from early universe phenomena (e.g., cosmic strings, superstrings, or inflation), and the dynamics of space-time itself (e.g., a coalescing black hole binary and its coalescence to form a single black hole). They are detectable by the small disturbance they cause in the light travel time between some light source and an observer. NANOGrav exploits radio pulsars as both the light (radio) source and the clock against which the light travel time is measured. In an array of radio pulsars gravitational waves manifest themselves as correlated disturbances in the pulse arrival times. The timing precision of today’s best measured pulsars is less than 100 ns. With improved instrumentation and signal-to-noise it is widely believed that the next decade could see a pulsar timing network of 100 pulsars each with better than 100 ns timing precision. Such a *pulsar timing array* (PTA), observed with a regular cadence of days to weeks, would be capable of observing supermassive black hole binaries following galactic mergers, relic radiation from early universe phenomena such as cosmic strings, cosmic superstrings, or inflation, and more generally providing a vantage on the universe whose revolutionary potential has not been seen in the 400 years since Galileo first turned a telescope to the heavens.

The parameters that determine the sensitivity of a pulsar timing array to gravitational waves are the number, timing precision, and sky distribution of the array’s best pulsars. To achieve its goals, NANOGrav requires improved pulsar timing, with higher cadence and more regular timing of the array pulsars; searches for new, stable millisecond pulsars to enlarge the array size (especially in the northern sky); and improved statistical analyses of timing residuals to increase the robustness of gravitational wave detection at lower signal strength.

To make its observations, NANOGrav can exploit existing and new radio astronomy telescope infrastructure. The broader bandwidth of modern pulsar back-end instrumentation presents the opportunity to dramatically increase pulsar timing precision. Taking advantage of broad-band observations requires the development of techniques that account for interstellar medium propagation effects, frequency dependent profile changes, polarization calibration errors, and radio frequency interference (RFI) mitigation. With the appropriate computational infrastructure, NANOGrav can take excellent advantage of the proposed phased-array radio telescopes (e.g., the Allen Telescope Array and the Square Kilometer Array) and, in turn, drive the development of a non-imaging timing mode. Such development will prepare the NANOGrav community to take full advantage of the next generation of radio astronomy hardware.

## 2 Key Science Goals

Gravitation powers and governs the most energetic processes in the cosmos, from supernovae, to gamma-ray bursts, to quasars, to the coalescence of supermassive black hole binaries following galaxy mergers. Gravitational waves associated with these phenomena are the only direct probe of the central engines that power them. Gravitational waves are also the most direct diagnostic of the structure of space-time, and of the processes taking place in the earliest moments in the life of the Universe. Microhertz or lower frequency gravitational waves associated with, e.g., coalescing supermassive black holes or black hole binaries cannot be detected using conventional ground- or space-based detectors. They can, however, be detected through the correlated imprint they leave on the timing residuals derived from observations of an array of pulsars. The combination of expected source strengths, timing precision and number of low-timing-noise pulsars suggest strongly that gravitational wave observations with a pulsar timing array will be possible within the next decade (Jenet et al. 2005). The North American Nanohertz Observatory for Gravitational Waves — NANOGrav — is a consortium of astronomers whose goal is to observe nHz– $\mu$ Hz frequency gravitational waves and use those observations as a tool of observational astronomy, contributing principally to

1. Understanding the co-evolution of galaxies and supermassive black holes;
2. Searching for signatures of early-universe or exotic physics processes (e.g., inflation or cosmic strings);
3. Probing the nature of space-time, including the search for quantum gravity corrections to classical gravity; and
4. Discovering sources of gravitational waves previously unrecognized.

### 2.1 Background

Pulsars are exquisite clocks. Among all pulsars, the class of millisecond pulsars are especially precise timekeepers. Owing their stability to their large moment of inertia and “seismic” stability, the rms timing residuals — the delays between the expected and actual arrival time of pulses — for several millisecond pulsars are below 100 ns. The timing residuals have steadily improved with improved instrumentation and signal-to-noise, strongly suggesting that we have not yet reached the intrinsic pulsar timing noise limits.

A pulsar’s regular tempo makes it an invaluable tool for exploring relativistic phenomena where precise and accurate measurement of duration or interval is critical. The Hulse & Taylor (1975) discovery of the pulsar PSR B1913+16 in a binary system enabled the precision measurement of multiple relativistic corrections to Newtonian gravity, culminating with a measurement of the orbital period decay rate with precision  $10^{-15}$  s/s (Weisberg & Taylor 2005). The agreement between this  $\dot{P}_b$  and general relativity’s prediction for the orbital period evolution owing to gravitational wave emission — currently 0.2% — was responsible for the first and, currently, best observational evidence for gravitational wave emission. The agreement between measured binary period decay and the predictions of general relativity provided by PSR B1913+16 and other, subsequently discovered binary systems is at this writing the *only* test of general relativity theory in its dynamical sector.

In a binary pulsar system gravitational wave emission is inferred from its dissipative effect on the binary’s orbital period. The pulsar is, in this case, the precision clock that enables us to characterize the evolving orbit. Gravitational waves also affect space-time: in particular, they disturb the time it takes a pulse to propagate between a pulsar and a radio telescope observatory. The passing wave’s effect on this propagation time depends on the wave propagation direction and polarization relative to the pulsar-observatory line-of-sight. Correlations among the timing residuals measured for an array of pulsars along different lines-of-sight can thus reveal the presence of gravitational waves, their propagation direction and their polarization.

The pulsar timing residual associated with gravitational waves crossing the observatory-pulsar line-of-sight is equal to geometrical factors times the time-integrated gravitational wave strain, commonly denoted  $h_{ij}$ , along the electromagnetic wave-front as it propagates from the pulsar to the observatory. *The ability to use the pulsar timing array to measure the equivalent of nanosecond or less timing residuals induced by gravitational waves, corresponding to a strain sensitivity to broadband bursts of order  $h \sim 6 \times 10^{-15} (f_{\text{gw}}/10^{-6} \text{ Hz})$ , is well within the capability of next-decade pulsar timing array measurements.* The subsections below describe how NANOGrav will exploit this sensitivity to achieve its principal science goals.

## 2.2 Understand the co-evolution of galaxies and supermassive black holes

Massive black holes appear to be ubiquitous inhabitants of all galactic nuclei with spheroids (Richstone et al. 1998). The properties of these black holes also appear to be strongly correlated with the properties of their hosts (Magorrian et al. 1998; Gebhardt et al. 2000; Ferrarese & Merritt 2000). Together these observations support a co-evolution scenario for massive black holes and galaxies that would extend to other phenomena (e.g., energy radiated by a quasar over its lifetime) previously regarded as independent (Hopkins et al. 2008). In this scenario, the correlations between nuclear black holes and their hosts are driven by the processes that follow from galaxy mergers. Following a merger, dynamical friction leads the massive black holes to sink to the center of the merger remnant. At the center they form a bound binary that loses its orbital energy through interactions with stars and gas. Eventually the stars and gas are exhausted and the coalescence is driven by gravitational wave emission (Begelman, Blandford, & Rees 1980).

The timing residual for a massive black hole binary is readily estimated. For a binary with total mass  $M/(1+z)$  at luminosity distance  $d_L$  radiating gravitational waves with observed frequency  $f$  the induced timing residual magnitude is, up to geometric factors,

$$\Delta\tau \sim 10 \text{ ns} \left( \frac{1 \text{ Gpc}}{d_L} \right) \left( \frac{M}{10^9 M_\odot} \right)^{5/3} \left( \frac{10^{-7} \text{ Hz}}{f} \right)^{1/3} \quad (1)$$

Correlated timing residuals of this magnitude are well within reach of pulsar timing array measurements in the next decade.

Observations with an array of 20 pulsars, each with 100 ns rms timing residual noise, are expected to observe a confusion-limited gravitational wave “background” signal in the

nHz– $\mu$ Hz window arising from binaries with masses  $> 10^8 M_\odot$  (Jaffe & Backer 2003; Jenet et al. 2005; Sesana, Vecchio, & Colacino 2008). On top of this confusion-limited signal should be a handful of individually resolvable binaries, of greater mass, at distances  $z \lesssim 2$  (Sesana, Vecchio, & Volonteri 2009). The level and spectral shape of the background signal reflects the abundance of massive binaries and their eccentricity, which are determined by the dynamical processes acting at sub-parsec scales. The background spectral shape is also strongly constrained by the still disputed shape of the high-mass end of the black hole mass function (e.g., Lauer et al. (2007)). Observations of the gravitational wave background in this band will thus provide valuable insights into the poorly known accretion and feedback processes that govern the growth of the massive black holes that play a dominant role in shaping the structure of giant galaxies and galaxy clusters.

The observation of individually resolvable binaries offers even more exciting prospects. For these systems the binary’s location on the sky, orbital plane inclination, component mass ratio and total mass to luminosity distance ratio can all be determined. The mass ratio will be especially valuable for shaping our understanding of galaxy mergers and galaxy formation processes. If an electromagnetic counterpart can also be identified the binary system’s component masses will be determined (from the mass/distance ratio) and a unique laboratory for studying accretion physics and the interplay between black holes and their host galaxies uncovered.

### 2.3 Search for the signatures of early-universe or exotic physics processes

Gravitational wave relics of early universe phenomena carry the signature of the physics at work in extreme environments otherwise inaccessible to observation or experiment. The relic radiation extends over wavelengths ranging from the present-day horizon scale (corresponding to  $3 \times 10^{-18}$  Hz) to a cut-off determined by the Planck scale and the expansion dynamics of the universe (which is expected to be of order gigahertz). Over this broad dynamic range in frequency, cosmic microwave background polarization measurements are sensitive to gravitational waves with wavelengths near the horizon scale, pulsar timing observations to radiation in the nHz– $\mu$ Hz band, the proposed Laser Interferometer Space Antenna (LISA) detector to the 0.1–100 mHz band, and the advanced Laser Interferometer Gravitational-Wave Observatory (LIGO) detector to the 10–1000 Hz band. While each of these different means of detecting relic gravitational waves thus explores a very different range of frequencies where the radiation spectrum is dominated by different processes, at present and into the foreseeable future pulsar timing experiments provide the tightest constraints on the parameter space of viable cosmic string and superstring models (Siemens, Mandic, & Creighton 2007).

In the nHz– $\mu$ Hz band relevant for pulsar timing array gravitational wave detection, the dominant early-universe contributions to the gravitational wave background are cosmic strings, cosmic superstrings, and inflation. Cosmic strings are one-dimensional topological defects associated with a symmetry-breaking phase transition (Kibble 1976) in the early universe. Originally proposed to provide the seeds of structure formation, by the late 1990’s this motivation evaporated when they were found to be inconsistent with the distribution of the large-scale structure and the observation of acoustic peaks in the cosmic microwave

background. Nevertheless, cosmic strings, which are a generic phenomenon of supersymmetric grand unified theories (Jeannerot, Rocher, & Sakellariadou 2003), remain viable as a sub-dominant contributor to the cosmic gravitational wave background (Wyman, Pogosian, & Wasserman 2005). *Superstring*-motivated inflation models can also lead to artifacts that behave like cosmic strings; such artifacts are called cosmic superstrings (Polchinski 2005). Both cosmic strings and superstrings support oscillations, cusps and kinks, which in turn give rise to a stochastic gravitational wave “background” and gravitational wave bursts (Damour & Vilenkin 2001, 2005). Finally, inflation gives rise to a stochastic cosmological gravitational wave background by powering a parametric amplification of gravitational wave zero-point quantum fluctuations (Grishchuk 2005).

Remarkably, at  $10^{-7}$  Hz current predictions for the contribution of cosmic strings, superstrings or inflation to the stochastic gravitational wave background are all of comparable magnitude ( $h_{\text{cs}} \sim 10^{-16}\text{--}10^{-14}$  and  $h_{\text{infl}} \sim 10^{-17}\text{--}10^{-15}$ ) and within the reach of pulsar timing observations in the next decade (Jenet et al. 2006). The spectral index of each of these contributions differ, offering the prospect that the contributions can be separately identified from the spectrum of the background measured over the nHz– $\mu$ Hz bandwidth. (These amplitudes are also comparable with that expected from the confusion limit of supermassive black hole binaries, which has yet a different spectral index allowing it to also be distinguished.)

## 2.4 Probe the nature of space-time

**Quantum gravity.** Astronomical observations involve the largest imaginable scales of distance and time. The leverage provided by the vast scales probed by astronomical observations enable tests of physical theories that cannot be approached in terrestrial laboratories or by conventional experiment. Classic examples include measurements bounding possible time-evolution of the fine structure “constant,” which involve comparison of spectral line multiplet frequencies made at present and when the universe was considerably younger (Darling 2003; Kozlov et al. 2004) and bounds on the mass of the photon (Adelberger, Dvali, & Gruzinov 2007; Chibisov 1976). In a similar fashion, gravitational wave observations of sources at cosmological distances will make possible tests of string and other theories that attempt to merge quantum mechanics and gravity.

All modern routes leading to a quantum theory of gravity —e.g., perturbative quantum gravitational one-loop exact correction to the global chiral current in the standard model, string theory, and loop quantum gravity — require modification of the classical Einstein-Hilbert action by the addition of a parity-violating Chern-Simons term. This quantum correction affects only the gravitation sector of these theories; it has no effect on electromagnetism or electromagnetic phenomena. While the correction may be intrinsically small, its effect on gravitational waves accumulates secularly as the wave propagates. Observation of gravitational waves that have propagated over cosmological distances will be capable of measuring or bounding even a small birefringence, providing evidence that will help to guide the unification of gravitation and quantum theories (Alexander, Finn, & Yunes 2008).

**General relativity, or beyond?** Solar system observations and equivalence principle experiments have shown that gravitation is almost certainly described by a metric theory (Will 2006). Einstein’s theory of gravity — general relativity — is the simplest such theory.

While general relativity has been extremely successful in describing gravitational physics there still remain many feasible alternative gravity theories. Interest in such theories has increased recently due to discoveries in galactic dynamics and cosmology, i.e., dark matter and cosmic acceleration (see the review (Sanders & McGaugh 2002)). General relativity makes an unambiguous prediction for the number of gravitational wave polarization states (two), which is fewer than allowed (six) by these viable alternative theories (Eardley et al. 1973; Eardley, Lee, & Lightman 1973). Thus, gravitational wave observations that can independently resolve the different polarization states enable a test that can distinguish between general relativity and other viable alternative theories of gravity.

Terrestrial and space-based gravitational wave detectors are each sensitive to a single linear combination of the polarization states associated with a passing gravitational wave. When data from multiple independent detectors are joined and analyzed coherently the ability to resolve the different polarization states is increased. For this purpose, each pulsar line-of-sight is an independent detector: i.e., the timing residuals associated with different pulsars each respond to a different linear combination of gravitational wave polarizations. Over the next decade there will likely be no more than three large independent terrestrial gravitational wave detectors (detectors at the two LIGO sites, and the French-Italian Virgo detector) and at most one space-based detector (LISA). There are already about 10 independent pulsar baselines with of order 100 ns timing noise and good reason to expect there to be more than 100 pulsar baselines with sub 100 ns timing noise by the end of the next decade. The size of this “detector array”, which cannot be matched by ground or space-based experiment, enables observations capable of distinguishing radiation modes associated with each polarization state and, thus, distinguishing between general relativity and other alternative theories of gravity (Lee, Jenet, & Price 2008).

## 2.5 Discover previously unanticipated gravitational wave sources

Nothing more exemplifies the transformative potential of gravitational wave astronomy than the possibility of discovering gravitational waves from unforeseen sources or processes. It is cliche to say that the history of astronomical observation is a catalog of surprises as new observational channels create new opportunities for discovery. Behind the cliche, however, is the truth from which it is made: the range of our direct experience, from which our knowledge and understanding of the natural world is formed, is very limited when compared to the range of scales and environments present in the universe. For this reason, great leaps in understanding follow from great leaps in observational capability. This decade the gravitational wave universe lays claim to the imagination as the great, unexplored frontier ripe for exploration. NANOGrav, its international partners, and other gravitational wave experiments, are the means that will carry us into this new frontier, from which we may enrich our understanding of the known, and the as yet unknown, cosmos.

### 3 Technical Overview

A single plane gravitational wave with a well defined frequency will cause the measured time-of-arrival of individual pulsar pulses to oscillate with time; most of our expected signals involve linear superpositions of the effects of many such waves. The NANOGrav goal is to detect such effects by careful examination of long-term pulsar timing signatures. **The single most important performance parameter is the root-mean-square (RMS) “timing residual”  $\sigma_i$  for each of the pulsars in our sample.** The fiducial goal of a detection of a stochastic background of gravitational waves requires roughly 5 years of timing 20 pulsars, each with  $\sigma_i$  of  $\sim 100$  ns (Jenet et al. 2005). Our achievable timing precision has been steadily improving over the last two decades (Demorest & Jenet 2009), as shown in Fig. 1. Given a concerted effort by NANOGrav and its international collaborators, we expect to continue this trend and achieve our primary goal — the detection of GWs in the nanohertz range — before the end of the decade. This will require a substantial amount of observing time on large radio telescopes, together with computational resources for data analysis and human resources for algorithm development.

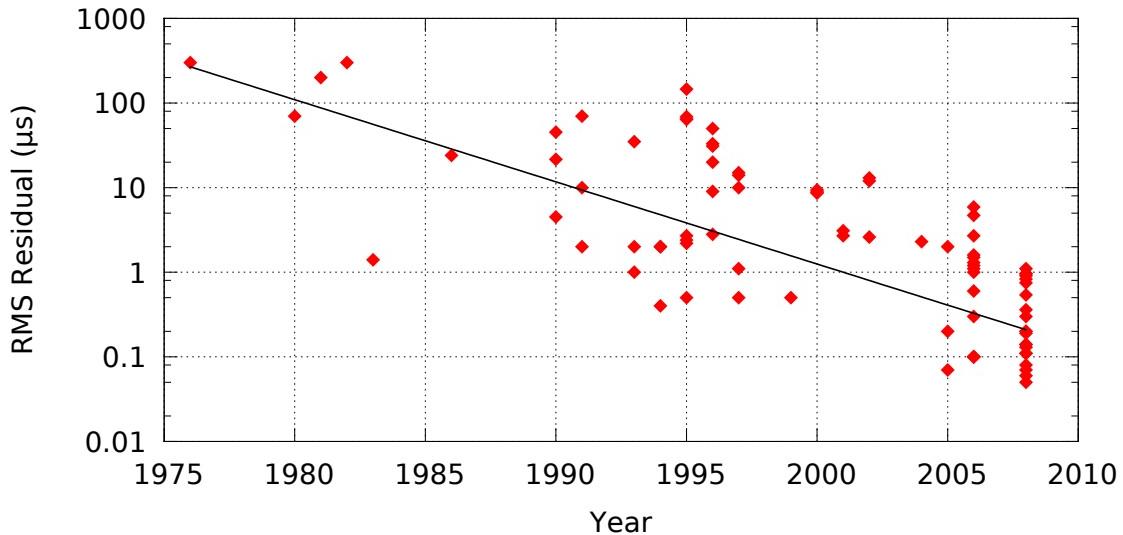


Figure 1: Published RMS pulsar timing residuals versus time, showing exponential improvement (Demorest & Jenet 2009). Continued improvement should allow us to detect GWs within the next decade.

In this section, the NANOGrav technical program is outlined in detail, including signal detection, infrastructure requirements and technology drivers. References to specific technology driver sections are indicated. Current pitfalls are described together with NANOGrav’s plans for mitigating these problems. The full NANOGrav process can be broken into 4 stages: observation, time-of-arrival (TOA) calculation, derivation of timing models and residuals for the individual pulsars, and combining pulsar residuals to detect or improve upper limits on the GWB. In addition to the technical issues discussed here, we have another crucial need: to recruit and train young scientists at the undergraduate, graduate and postdoctoral levels.

### 3.1 Observations

The discovery of millisecond pulsars (MSPs) (Backer et al. 1982) revolutionized high-precision timing. MSP rotational phases can often be determined within a few  $\mu\text{s}$  or better in a single observation. The sensitivity of a pulsar timing array (PTA) is determined by the number and distribution of the pulsars in the array, the cadence with which they are observed, and the precision with which the pulse times of arrival are measured. Currently there are several pulsars for which the RMS timing residual  $\sigma_i$  approaches 100 ns and roughly 20 more with residuals less than 1  $\mu\text{s}$ .

**Accomplishing our goal of characterizing GW sources requires routine access to significant amounts of time on large radio telescopes, plus observing (backend) instrumentation capable of removing the worst propagation effects and accumulating the pulsar signal over wide bandwidths. Specifically, it translates into approximately 10% of the observing time on the Arecibo telescope and about 20% of the time on the Green Bank Telescope (see §3.6).**

We need large telescopes for this effort since the pulsar signal strength increases directly with telescope collecting area, with a corresponding effect on the GW signal. We also need large integration times with large bandwidths since the signal strength is proportional to the square root of the product of the observing bandwidth and integration time. Finally, frequent observing epochs are necessary: first to ameliorate interstellar propagation effects and then to prevent long gaps that risk introducing spurious effects into the timing residuals used for the GW analysis. Extreme equipment stability and access to the highest quality clock information are paramount during the observation and between observation epochs. Comparison of observational data between observatories and simultaneous observations with multiple telescopes is also crucial.

Our observation requirements are also dictated by the need to correct for the propagation of the radio pulsar signal through the interstellar medium (ISM). This has two major effects on the observed signal. The first, dispersive smearing, can be very well modeled as a frequency-dependent filter acting on the data stream, and, hence, be accurately removed by applying the inverse filter to the pre-detection data. This coherent dedispersion technique (Hankins & Rickett 1975), which is computationally intensive, is currently in use across bandwidths of the order of 100 MHz at many of the world’s largest telescopes. In 2009, it will be implemented across many hundreds of MHz in the GUPPI<sup>1</sup> at the GBT, meaning that fully dedispersed pulsar signals will be available simultaneously across any one of the GBT’s sub-3-GHz receivers. Similar backends are required for Arecibo (the newly installed Mock spectrometers have wide bandwidth but cannot perform coherent dedispersion) and for future observatories.

A far more challenging problem is multipath propagation of the pulsar signal in the ISM, which causes scintillation of the signal via constructive and destructive interference at the Earth, and also significant scattering for many pulsars, especially at low frequencies. Substantial efforts are underway to model the multipath effects and find ways to eliminate them from the pulse profiles (§4.2). Both the multipath scattering and dispersive effects are observed to change on timescales as short as one day and as long as decades (e.g. Kaspi, Taylor, & Ryba 1994; Cordes et al. 1990; Hemberger & Stinebring 2008). This makes multifrequency

---

<sup>1</sup>Green Bank Ultimate Pulsar Processing Instrument

(at least two receivers, each with wide bandwidths, such that at least 4 separate frequency bands are accessible) observations a requirement at each epoch.

An ever-increasing challenge is the presence of radio-frequency interference (RFI) in the signal. Our group has demonstrated initial success with broadband and narrowband software RFI identification and excision before and during coherent dedispersion (Stairs et al. 2000), but we need efficient versions of these algorithms in order to implement them routinely in the wideband real-time instruments. Achieving this goal and exploring alternate RFI mitigation schemes based on FPGA technology (Kesteven et al. 2005) are some of our priorities for the next few years (§4.1). We will also determine the exact effect of such techniques on the final recorded signal and TOAs by performing an end-to-end simulation of the pulsar signal path encompassing everything from propagation effects to calibration corrections.

## 3.2 Pulse Time of Arrival Determination

For a given pulsar, the raw data acquired during an observing epoch consists of a set of profiles, coherently dedispersed within several channels, detected as self- and cross-polarization products, folded synchronously using the best available ephemeris, and saved to disk roughly every minute.

Turning these recorded profiles into times of arrival (TOAs) is a fundamental component of our analysis. In brief, it involves the cross-correlation of the epoch profile(s) with a standard template for the pulsar, and adding the resulting time offset to the recorded clock time. We have identified many facets of our current procedures which require significant development. **In most cases we know what the starting points for the improvements should be, and we have actually begun to implement them, but the person-power is currently lacking to accomplish them all effectively and on the necessary timescale.** We note that we are developing two largely independent software pipelines to facilitate error recognition and correction. Key items of concern are mentioned below and also highlighted in §4.1.

- We are implementing multiple techniques for computing the receiver parameters and, and we are comparing these methods.
- Standard profile determination is vitally important, as it represents our best guess at the true envelope of the radio emission from the pulsar, and we are often working to a precision that is  $\sim 1\%$  of our sampling interval.
- The TOA determination algorithm most used for the last 15 years has been a frequency-domain cross-correlation of the total-intensity profile with the standard template (Taylor 1992). Although this has worked fairly well, we plan to incorporate full-Stokes information and are considering Bayesian approaches to this problem.

## 3.3 Timing Model Analysis

Timing analysis consists of determining the time delay differences (or residuals)  $R(t)$  between the observed TOAs and those predicted by a multi-parameter timing model. Parameters going into the timing model always include: information about clock errors; the Earth's

position and motion in the solar system; the pulsar’s rotational frequency and the first derivative of that frequency; the position, proper motion, and parallax of the pulsar; and information about radio frequency delays due to propagation through the ionized interstellar medium. In the common case that the MSP is in orbit around another compact object, full information about the binary orbit – including post-Keplerian parameters that, in turn, are a unique probe of general relativity – is also needed to specify the timing model. Development of this timing model for a pulsar is done in a boot-strap fashion over months to years of observation and analysis.

### 3.4 Correlation Analysis

Because our observing timescales are on the order of years, we will be sensitive to GWs with frequencies in the nanohertz regime. Such waves would most likely be due to coalescing supermassive binary black holes. When the timing analysis is applied to such a signal the dominant term in the residuals would typically be a cubic polynomial that grows in amplitude and changes in character as the time span of the observations increases. In the simplest case of a sine wave and a large GW signal from a nearby binary, it may be possible to detect this signal in the residuals of a single pulsar. However, the PTA forms a more robust and sensitive detector by allowing searches for correlations between the residuals from many pulsars distributed across the sky. The characteristic gravitational strain to which we are sensitive scales as  $h_c \propto \sigma/(T^2\sqrt{N})$ , where  $\sigma$  is the rms timing residual,  $N$  is the number of pulsars in the array, and  $T$  is the total time span of the observations (Jenet et al. 2005; Kaspi, Taylor, & Ryba 1994).

In addition to the obvious sensitivity gained from a larger number of pulsars, the additional importance of the PTA follows from its ability to distinguish between a gravitational wave signal and the many forms of noise – clock errors, solar system ephemeris errors, ISM propagation noise, timing noise intrinsic to pulsars, etc. – that will be present in the data. The unique character of an isotropic background of gravitational waves leads to a well defined expected correlation between the residuals that only depends on the angular separation angle between the pulsars in the array (Hellings & Downs, 1983; Jenet et al., 2005). Individual gravitational waves will also have a unique correlation signature but with a preferred direction on the sky. By contrast, an error in the reference clock will affect all the pulsars the same way (a monopolar signature), and a position or velocity error in the solar system ephemerides will cause an error that has a preferred direction with a dipolar signature around the sky. Several sources of error, e.g. rotational instabilities of individual pulsars or propagation delays along particular lines of sight, will lead to uncorrelated errors between pulsars.

Due to the necessity of determining the timing model parameters from the pulsar data itself, the correlation analysis will be insensitive to GWs at particular frequencies (Blandford, Romani, & Narayan 1984). For example, the presence of semiannual (parallax) and annual (positional) terms as well as the need to fit for the pulsar period and period derivative removes power from the GW signal with periods of 6 months and a year as well as removing a linear and quadratic term from the rotational phase as a function of time. In general, frequentist methods have been used to search the post-fit timing residuals for the correlated signal expected from a stochastic GW background (Jenet et al. 2005) and have set the most

stringent limits to date on the GW background. However, accounting for the GW power absorbed by the timing model fit has often been problematic for these algorithms. We plan to explore Bayesian approaches to detecting and placing limits on the GW background that will more naturally and correctly account for these effects (McHugh et al. 1996; Anholm et al. 2008; van Haasteren et al. 2008). Fully developing such methods is a cornerstone of our project, since they will lead to more robust upper limits, a better understanding of the errors in our data, and ultimately a detection of GWs within a shorter time span.

### 3.5 New Pulsars for the Timing Array

To fully characterize the gravitational waveforms and to maximize the scientific return, more precision MSPs are needed. Finding pulsars in directions widely separated from the current set of objects is important since the correlation over a wide range of separations is the key to detecting a GW background signal. While the actual performance of a PTA will depend upon the specific properties of the pulsars in the PTA, simple test cases suggest that the significance of a GW detection  $S_{\text{GW}}$  in an  $N$ -pulsar PTA scales roughly as  $S_{\text{GW}} \propto N$ , provided that all of the pulsars being timed have a sub-microsecond TOA precision.

Three current searches are now turning up such objects: the PALFA L-band multibeam survey at Arecibo, the GBT 350 MHz drift scan survey, and the new Parkes L-band Digital Survey. A new GBT low-frequency pulsar survey would be particularly advantageous because it would increase the number of Northern hemisphere pulsars. This area of the sky is currently under-represented in MSP targets. Since pulsar flux increases at low radio frequencies (from  $\sim 1$  mJy at 1.4 GHz to  $\sim 10$  mJy at 400 MHz), this survey will identify many new pulsars. We are also considering a 1.4 GHz GBT survey using a multibeam receiver.

In all of these searches, advances in computational hardware mean that more complicated signal processing can be undertaken. Many millisecond pulsars are in binaries, a reflection of their origin. However, finding pulsars in binaries often requires at least searching for accelerated pulsed signals, if not a search over binary parameters, in addition to the standard pulse period and dispersion searches. Objects in very tight orbits or a pulsar orbiting a black hole would, of course, produce spectacular collateral science to the main NANOGrav effort.

### 3.6 Justification of Required Telescope Time

In our Science white paper (Demorest et al 2009) we summarized observing time dedicated to this project in terms of 100-m telescope equivalent time. World resources currently provide about 300 100-m hours per month. We estimate that at minimum, GW detection requires  $\sim 500$  100-m hours per month, based on observing 20 pulsars every 2 weeks, for 3 hours at each of 4 radio frequencies in order to obtain 100-ns or better timing precision. To begin to characterize the GW sources requires at least twice as many pulsars (Lee, Jenet, & Price 2008) and so 1000 100-m hours of time each month. Upon purchasing 20% of the time on the GBT and 10% of the time at AO, we would have roughly 1300 100-m hours per month and would be able to characterize the stochastic GWB as well as continuous and burst sources. These numbers appear in our cost estimates.

## 4 Technology Drivers

### 4.1 Algorithm Development

Most of the stages in our technical program (§3) require improvement of existing algorithms or development of new ones. We argue that the entire class of algorithms developed through the NANOGrav activity (often in concert with our international collaborators) forms a technological tool kit that will enable science not just with the world's existing large radio telescopes but also with future telescopes such as the Square Kilometre Array. Here we recapitulate the most important algorithms that require progress.

- Radio-frequency interference (RFI) excision (§3.1), efficiently implementing existing software algorithms (Stairs et al. 2000), and exploring hardware-based solutions (Kesteven et al. 2005).
- Characterizing and reducing the effects of interstellar scattering (Hemberger & Stinebring 2008) and variable dispersion measures (Kaspi, Taylor, & Ryba 1994; You et al. 2007) (§3.1). As a relevant and currently feasible hardware development, we strongly advocate the building of receivers for the GBT and Arecibo that allow truly simultaneous observations at two widely spaced frequencies.
- Pulse profile calibration (§3.2) including full instrumental characterizations (Britton 2000; Johnston 2002; van Straten 2006).
- Standard-profile determination (§3.2) including frequency dependence (Lommen 2001).
- TOA computational methods (§3.2) should include frequency-evolving standard profiles (Lommen 2001) and take advantage of full-Stokes information (van Straten 2006) and/or Bayesian methods.
- Pulsar timing algorithms are under continuous international development (Edwards, Hobbs, & Manchester 2006) (§3.3).
- End-to-end simulation of the pulsar signal data path, including ISM effects, telescope reception, and our full analysis pipeline (§3.1)
- GWB analysis algorithms (§ 3.4) are still in an early stage of development, with both frequentist and Bayesian techniques being explored (Jenet et al. 2006; Anholm et al. 2008; van Haasteren et al. 2008).

## 4.2 Removing the Effects of the Interstellar Medium

Although we have made excellent progress in reducing rms residuals for many sources, we know that time-dependent propagation delays present a serious challenge in attaining 100 ns rms residuals toward many pulsars. We single this out as a technology driver because an extensive R&D program will be required, and we are just beginning to recognize the severity of the problem. As an example, we show in Fig. 2 the results of 270 days of monitoring the inferred time delay of the signal from one pulsar. At radio frequencies around 1400 MHz, a commonly used frequency for many of our high precision timing data, we see the inferred time delay varying from about  $0.2 - 2 \mu\text{s}$  with several abrupt fluctuations.

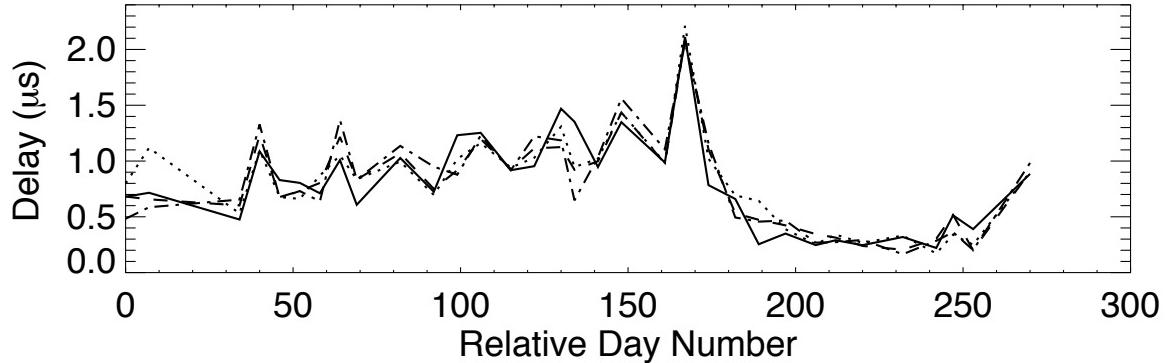


Figure 2: Time delay through the ionized interstellar medium toward the pulsar PSR B1737+13, estimated using spectral methods at four separate frequencies near 1400 MHz (adapted from Hemberger and Stinebring 2008).

In an inhomogeneous plasma like the interstellar medium, the majority of rays from the pulsar travel along a cigar-shaped volume tapered at both ends. The angular beam waist (scattering radius) is highly wavelength dependent ( $\theta_s \propto \lambda^{2.2}$ ) and varies with time. Furthermore, this description of the ray paths (including the exponent of 2.2) depends in a complicated fashion on the inhomogeneity spectrum and the distribution of material along the line of sight (Cordes & Rickett 1998), none of which is known *a priori*.

We have considerable expertise on this problem in our group, and several of our international collaborators are working hard on the problem, too. One straightforward approach would be to use spectral information like that underlying the results in Fig. 2 and simply subtract the scattering time delay from the TOA. This requires excellent signal-to-noise ratio, achievable for only the strongest pulsars and the largest telescopes, and we have not yet demonstrated that it effectively reduces the rms residuals. A more ambitious approach, that may be necessary in some cases, is to use pre-detection information to fully model the phase front of the wave as it passes through the interstellar medium. We have taken some initial steps in this direction (e.g. Walker and Stinebring 2005). This approach can be thought of as coherent descattering. Unlike coherent dedispersion, however, this correction requires an iterative modeling approach that, again, requires high sensitivity and frequent (daily) sampling of the data. Although the challenges we face in this area are daunting, we note that coherent dedispersion — one of the current bedrocks of our observing procedure — was initially considered farfetched by many, and it has proven to be a highly effective and essential tool for precision pulsar timing.

### 4.3 Pulsar Timing with Interferometers

Over the next decade, a number of new instruments will come on-line, including the EVLA, the Australian SKA Pathfinder (ASKAP), and potentially an expanded Allen Telescope Array (ATA-256). Looking to the latter half of the next decade and beyond, the Square Kilometre Array (SKA) is expected to deliver even larger improvements in sensitivity. All of these new telescopes are interferometers, in contrast to the current workhorse telescopes (Arecibo, GBT, and Parkes).

In addition to the possibility of simply additional sensitivity,<sup>2</sup> interferometers offer the possibility of improved scheduling and efficiency. In principle, via *sub-arrays*, an array can be split so that multiple pulsars can be timed simultaneously. Implementing such a timing program, and achieving an optimum schedule, will require careful attention to the flux densities of the pulsars in the PTA and the range of frequencies needed for each pulsar (to compensate for interstellar scattering effects).

Further, the standard operational mode of interferometers is for imaging, while pulsar timing represents a *non-imaging* operational mode. Rather than visibilities, a *phased-array* mode must be implemented in (or in addition to) the correlator. While conceptually a phased-array mode is not difficult, because pulsar observing has not been conducted traditionally with interferometers, the implementation of pulsar-specific processing typically lags that of standard imaging processing. As specific examples, implementation of pulsar modes in the VLBA correlator required approximately 5 years after the formal dedication of the telescope, and the initial operational modes of the EVLA correlator will not include pulsar modes.

There is some modest experience with using an interferometer as part of a PTA, being developed primarily through the EPTA and its use of the Westerbork Synthesis Radio Telescope (WSRT). Even this experience, however, shows that careful attention must be paid to the process of phasing the array and extracting TOAs from the resulting phased-array data.

---

<sup>2</sup> The EVLA has an effective aperture equivalent to that of a 130-m single-dish telescope.

## 5 Activity Organization

**NANOGrav Internal Organization.** As detailed in §3, the requirements of this project include many aspects ranging from observing to algorithm development to data analysis, and therefore necessitate a large collaboration. At present, NANOGrav consists of 25 researchers spread primarily over a host of universities and colleges across the United States and Canada, serving an ethnically and socio-economically diverse population of undergraduate and graduate students. Our specialities include pulsar searching (Cordes, Freire, Kaspi, Lorimer, McLaughlin, Nice, Ransom, Stairs), timing (Backer, Demorest, Ferdman, Freire, Gonzalez, Jenet, Kaspi, Lommen, Lorimer, Nice, Ransom, Stairs, Verbiest), interstellar propagation (Cordes, Stinebring), gravitational wave detection (Jenet, Lommen), instrumentation (Demorest, Ransom, Stairs) and polarization (Demorest, Ferdman, Gonzalez, Rankin). All of these aspects are absolutely essential to meet our goals. We communicate with each other through weekly to monthly teleconferences, and collaborative visits. Basic sharing of data within NANOGrav has been implemented and we are moving towards common data formats and shared processing algorithms. By sharing algorithms instead of imposing a single standard, we ensure a redundancy that allows prompt identification and mitigation of conceptual and technical errors.

At present, NANOGrav activities are funded through the individual grants to US and Canadian faculty members as well as various fellowships awarded to some of the more junior members. The US members of NANOGrav recently applied for funds to further strengthen this collaboration through several measures. One of these is to increase the number of postdocs and students (both graduates and undergraduates), since 17 of the 25 current collaborators are senior researchers. We also plan extensive outreach efforts to involve high-school students in our research and pave the way for a new generation of gravitational wave astronomers. We aim to institute a yearly, week-long winter workshop dedicated to algorithm comparison, data analysis and resolving of data-related issues. Finally, dedicated computer hardware will increase the reliability and ease-of-use of our data sharing infrastructure.

**International Partnerships:** We organized the first international pulsar timing array (IPTA) meeting at the Arecibo observatory in August 2008. This event led to a formal agreement on collaboration between NANOGrav and its partner organizations in Europe (the EPTA) and Australia (the Parkes pulsar timing array or PPTA). We currently participate in monthly telecons with those groups and an agreement on data-sharing has been drafted that will soon enable regular combination of international timing data. This will allow rapid identification of technical failures or instabilities at any telescope and comparison and evaluation of algorithms used at different institutions. We also plan to involve colleagues in India since the Giant Meterwave Radio Telescope (GMRT) offers low-frequency capability that will greatly improve our achievable timing precision.

Along with the PPTA and EPTA, NANOGrav has recently joined the gravitational wave international committee (GWIC) and has strengthened its relation with LIGO (two of our current members are also in the LIGO collaboration). This wider collaboration is needed to fully benefit from the scientific complementarity of the different GW detection efforts.

## 6 Activity Schedule

**Timing Observations** As our sensitivity to a gravitational wave background strongly depends on the length and cadence of the timing data sets, regular and frequent timing observations will be ongoing throughout the lifetime of the project, at all available telescopes.

**Source Selection and Pulsar Searches.** High precision timing is crucial to achieving our scientific goals. Since the achievable timing precision varies from pulsar to pulsar, state-of-the-art surveys are needed to find any rapidly rotating MSPs to which previous surveys may not have been sensitive. Planned and currently ongoing searches are:

- **GBT drift-scan survey:** Observing has finished, processing is expected to end in 2010.
- **Arecibo Northern Galactic plane survey:** Observations are underway. The survey and data processing are expected to be finished by 2012.
- **GBT all-Northern-sky low-frequency survey:** This survey is intended to start after the processing of the drift-scan survey is complete. The survey and analysis should take place from 2010 to 2014.
- **Arecibo mid-latitude survey:** This survey would follow the Northern Galactic plane survey and run from 2012 to 2015.

MSPs detected in any of these surveys may help us to achieve our scientific goals, but the timing potential of these new pulsars must be verified. Continuous monitoring will therefore have to be undertaken for at least a year or more upon discovery. This monitoring and source selection will hence be ongoing until 2017.

**Algorithm Development** Four separate development efforts are planned. Details on these activities are provided in §3 and §4.1.

- Basic calibration and timing tools described in §3.2. We anticipate full integration of these algorithms into the current, data processing pipeline by the end of 2012, allowing these new techniques to be used in the analysis of the ongoing timing observations from the beginning of the decade onwards.
- Characterization of the effects of the ISM and the implementation of algorithms to mitigate such effects is a substantial effort that is expected not to be completed until late 2013. The second activity in at this stage is the creation of an end-to-end simulation of the

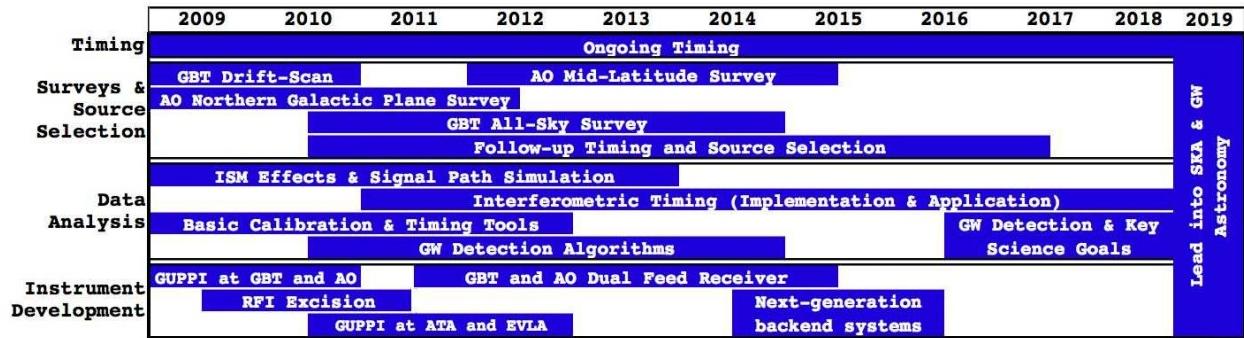


Figure 3: Timeline for the planned NANOGrav activities. A full description is given in the text.

signal data path will be constructed for system-analytic purposes. The construction of this simulation and the interpretation of its results are expected to continue into 2013 as well.

- GW detection algorithms based on both frequentist and Bayesian principles (see also §4.1), will be finalized by 2015.
- As outlined in §4.3, interferometers will become increasingly important in pulsar timing. Optimization of our techniques and algorithms to deal with interferometric telescopes will be completed by the latter half of the decade.

**Instrument Development** Over the course of the decade, we expect to make substantial improvements in pulsar instrumentation, in three phases:

- **Current backend systems:** The GUPPI backend, currently under development at Green Bank, will provide high dynamic range, full polarization, and coherent dedispersion capabilities over an order of magnitude larger bandwidth than previously possible. By the end of 2009, the full functionality of GUPPI will be implemented. We plan to install a copy of GUPPI at Arecibo by the end of 2010, and at the ATA and EVLA by 2012. Concurrently we will investigate the possibility of adding real-time RFI excision based on established techniques (Stairs et al. 2000) to these systems.
- **Receiver systems:** We will encourage dual-frequency receivers to be installed at the GBT and Arecibo by 2015. These receivers will enable simultaneous timing observations at high and low frequencies. In addition to improving observing efficiency, this will allow us to better quantify and correct time-variable ISM effects. We will also investigate the feasibility of a multibeam receiver for the GBT, which would dramatically improve the speed and sensitivity of future pulsar surveys.
- **Next-generation backend systems:** We plan to continue research and development of advanced next-generation backend systems, with a goal of upgrading the current (GUPPI-based) systems by 2015. The new systems will be designed to make use of the full bandwidth provided by the new dual-frequency receivers. They will also include advanced real-time RFI mitigation capability, and the ability to deal with specialized interferometer configurations such as subarrays or multiple beams.

**Activity Lifetime** Contemporaneously with advanced LIGO we expect the PTA sensitivity can reach a level where we will detect gravitational waves from supermassive black hole binaries, and detect or place astrophysically significant constraints on the gravitational signature of early universe phenomena (cf. §2.3). With these observations, we will introduce new tests of dynamical gravity (cf. §2.4) and explore black hole/galaxy co-evolution (cf. §2.2). The end of the 2010 decade also matches the projected beginning of the SKA era; our activities, and those of our international counterparts, will merge smoothly into intensive timing programs with the SKA and the commencement of GW astronomy as a science in itself. The efforts invested now into algorithm development and building up long timing baselines will ensure that the SKA can produce superb-quality science from the beginning.

## 7 Cost Estimates

In this section, we estimate the cost of running the NANOGrav activity itself. The operating costs include direct funding to the scientific community as well as facility costs.

Item	One-time Direct Funding	Direct Funding Cost/year	Facilities Cost/year	Decade total
AO Telescope Time	-	-	1M	10M
GBT Telescope Time	-	-	2M	20M
Personnel	-	2M	-	20M
Computing Hardware	300K	100K	-	1.3M
Travel and Communication	-	1M	-	10M
Hosting Meetings	-	300K	-	3M
GBT/AO receiver Upgrade	1M	-	-	1M
Recording Hardware Upgrade	300K	-	-	300K
Next Generation Backend Development	400K	-	-	400K
Total Cost				66M

Although we estimate that current levels of time allocated to us from the Green Bank and Arecibo Telescope will allow a detection of gravitational waves, more time will be needed to actually characterize gravitational wave sources. To achieve the scientific goals we have described in this document we estimate we need 10% of the observing time at the Arecibo observatory and 20% at the Green Bank radio observatory. The value of this is estimated to be about \$2 and \$1 million per year respectively for a total of \$3 million per year.

Aside from the facilities costs, the major operating cost of NANOGrav will be in supporting graduate students and a post-doctoral work force. (We already have the senior staff to supervise young scientists well in a variety of research and educational settings.) It is expected that this will be done through some combination of individual investigator grants together with a block grant to establish a “NANOGrav Institute.” We estimate \$2M per year to maintain 8 post-doctoral fellows, 10 graduate students, 2 hardware engineers, 2 computer scientists and summer salary for 10 faculty. This includes salary, overhead and benefits. The two Canadian groups also contribute personnel via their individual grants and some fellowships to their group members.

In addition to the personnel costs described above, we require funding for computers, travel, communication, and meetings; all oriented around organizing and maintaining international cooperation.

The installation of the GUPPI data taking systems at AO, ATA and EVLA will cost about \$100K each for hardware, for a total of \$300K. Installation costs are included in the engineering salaries above. Dual feed receivers for both AO and GBT will cost \$500K each for hardware. Next generation backend development will cost \$100K at each of the 4 telescopes (GBT, AO, ATA, and EVLA) for a total of \$400K.

Over the entire ten years, the estimated cost to the US of running the NANOGrav activity will be approximately \$66M (inclusive of the fractional operating costs of the national facilities used by this project). Considering the cost of other gravitational wave detection efforts, the NANOGrav activity is remarkably inexpensive.

## References

- Adelberger, E., Dvali, G., & Gruzinov, A. 2007, Phys. Rev. Lett., 98
- Alexander, S., Finn, L. S., & Yunes, N. 2008, Phys. Rev. D, 78, 066005–+
- Anholm, M., Ballmer, S., Creighton, J. D. E., Price, L. R., & Siemens, X. 2008, ArXiv e-prints, (astro-ph/0809.0701)
- Backer, D. C., Kulkarni, S. R., Heiles, C., Davis, M. M., & Goss, W. M. 1982, Nature, 300, 615–618
- Begelman, M. C., Blandford, R. D., & Rees, M. J. 1980, Nature, 287, 307–309
- Britton, M. C. 2000, ApJ, 532, 1240–1244
- Chibisov, G. 1976, Uspekhi Fizicheskikh Nauk, 119, 551–555
- Cordes, J. M., & Rickett, B. J. 1998, ApJ, 507, 846–860
- Cordes, J. M., Wolszczan, A., Dewey, R. J., Blaskiewicz, M., & Stinebring, D. R. 1990, ApJ, 349, 245
- Damour, T., & Vilenkin, A. 2001, Phys. Rev. D, 64, 064008–+
- Damour, T., & Vilenkin, A. 2005, Phys. Rev. D, 71, 063510–+
- Darling, J. 2003, Phys. Rev. Lett., 91
- Demorest, P., & Jenet, F. 2009, in preparation
- Eardley, D. M., Lee, D. L., & Lightman, A. P. 1973, Phys. Rev. D, 8, 3308–3321
- Eardley, D. M., Lee, D. L., Lightman, A. P., Wagoner, R. V., & Will, C. M. 1973, Phys. Rev. Lett., 30, 884–886
- Edwards, R. T., Hobbs, G. B., & Manchester, R. N. 2006, MNRAS, 372, 1549–1574
- Ferrarese, L., & Merritt, D. 2000, ApJL, 539, L9–L12
- Gebhardt, K., Kormendy, J., Ho, L. C., Bender, R., Bower, G., Dressler, A., Faber, S. M., Filippenko, A. V., Green, R., Grillmair, C., Lauer, T. R., Magorrian, J., Pinkney, J., Richstone, D., & Tremaine, S. 2000, ApJL, 543, L5–L8
- Grishchuk, L. P. 2005, Phys-Usp+, 48, 1235–1247
- Hankins, T. H., & Rickett, B. J. 1975, In Methods in Computational Physics Volume 14 — Radio Astronomy, New York: Academic Press, pp. 55–129
- Hemberger, D. A., & Stinebring, D. R. 2008, ApJL, 674, L37–L40
- Hopkins, P. F., Hernquist, L., Cox, T. J., & Kereš, D. 2008, ApJS, 175, 356–389

- Hulse, R. A., & Taylor, J. H. 1975, ApJL, 195, L51–L53
- Jaffe, A. H., & Backer, D. C. 2003, ApJ, 583, 616
- Jeannerot, R., Rocher, J., & Sakellariadou, M. 2003, Phys. Rev. D, 68, 103514
- Janet, F., Hobbs, G., Lee, K., & Manchester, R. 2005, ApJ, 625, L123–L126
- Janet, F. A., Hobbs, G. B., van Straten, W., Manchester, R. N., Bailes, M., Verbiest, J. P. W., Edwards, R. T., Hotan, A. W., Sarkissian, J. M., & Ord, S. M. 2006, ApJ, 653, 1571–1576
- Johnston, S. 2002, Pub. of the ASA, 19, 277–281
- Kaspi, V. M., Taylor, J. H., & Ryba, M. 1994, ApJ, 428, 713–728
- Kesteven, M., Hobbs, G., Clement, R., Dawson, B., Manchester, R., & Uppal, T. 2005, Radio Science, 40, 5–+
- Kibble, T. W. B. 1976, J. Phys., A9, 1387–1398
- Kozlov, M., Korol, V., Berengut, J., Dzuba, V., & Flambaum, V. 2004, Phys. Rev. A, 70
- Lauer, T. R., Tremaine, S., Richstone, D., & Faber, S. M. 2007, ApJ, 670, 249–260
- Lee, K. J., Janet, F. A., & Price, R. H. 2008, ApJ, 685, 1304–1319
- Lommen, A. N. 2001. Precision Multi-Telescope Timing of Millisecond Pulsars: New Limits on the Gravitational Wave Background and Other Results from the Pulsar Timing Array. Ph.D. thesis, University of California, Berkeley, Berkeley, CA, U.S.A.
- Magorrian, J., Tremaine, S., Richstone, D., Bender, R., Bower, G., Dressler, A., Faber, S. M., Gebhardt, K., Green, R., Grillmair, C., Kormendy, J., & Lauer, T. 1998, AJ, 115, 2285–2305
- Polchinski, J. 2005, AIP Conf. Proc., 743, 331–340
- Richstone, D., Ajhar, E. A., Bender, R., Bower, G., Dressler, A., Faber, S. M., Filippenko, A. V., Gebhardt, K., Green, R., Ho, L. C., Kormendy, J., Lauer, T. R., Magorrian, J., & Tremaine, S. 1998, Nature, 395, A14
- Sanders, R., & McGaugh, S. 2002, Annual Review of Astronomy and Astrophysics, 40, 263–317
- Sesana, A., Vecchio, A., & Colacino, C. N. 2008, MNRAS, 390, 192–209, (astro-ph/0804.4476)
- Sesana, A., Vecchio, A., & Volonteri, M. 2009, Gravitational waves from resolvable massive black hole binary systems and observations with Pulsar Timing Arrays, (astro-ph/0809.3412)
- Siemens, X., Mandic, V., & Creighton, J. 2007, Phys. Rev. Lett., 98, 111101

- Stairs, I. H., Splaver, E. M., Thorsett, S. E., Nice, D. J., & Taylor, J. H. 2000, MNRAS, 314, 459–467, (astro-ph/9912272)
- Taylor, J. H. 1992, Phil. Trans. Roy. Soc. Lond., 341, 117–134
- van Haasteren, R., Levin, Y., McDonald, P., & Lu, T. 2008, ArXiv e-prints, (astro-ph/0809.0791)
- van Straten, W. 2006, ApJ, 642, 1004–1011
- Walker, M. A., & Stinebring, D. R. 2005, MNRAS, 362, 1279–1285
- Weisberg, J. M., & Taylor, J. H. 2005, In Binary Radio Pulsars, F. Rasio and I. H. Stairs, eds., San Francisco: Astronomical Society of the Pacific, pp. 25–31
- Will, C. M. 2006, Living Reviews in Relativity, 9
- Wyman, M., Pogosian, L., & Wasserman, I. 2005, Phys. Rev. D, 72
- You, X.-P., Hobbs, G., Coles, W., Manchester, R. N., Edwards, R., Bailes, M., Sarkissian, J., Verbiest, J. P. W., van Straten, W., Hotan, A., Ord, S., Jenet, F., Bhat, N. D. R., & Teoh, A. 2007, MNRAS, 378, 493–506